On the cognitive argument for cost-benefit analysis

Christiansen, Andreas

Published in:
Ethical Theory and Moral Practice

DOI:
10.1007/s10677-018-9883-9

Publication date:
2018

Document version
Peer reviewed version

On the Cognitive Argument for Cost-Benefit Analysis

Andreas Christiansen

This paper has been accepted for publication in *Ethical Theory and Moral Practice*

Abstract

In a number of writings, Cass Sunstein has argued that we should use cost-benefit analysis as our primary approach to risk management, because cost-benefit analysis corrects for the cognitive biases that mar our thinking about risk. The paper critically evaluates this 'cognitive argument for cost-benefit analysis' and finds it wanting. Once we make distinctions between different cognitive errors and between different aspects of cost-benefit analysis, it becomes apparent that there are really two cognitive arguments, neither of which is successful as arguments for cost-benefit analysis as a whole. One argument shows that the analysis aspect of cost-benefit analysis is warranted because it corrects for false beliefs about the magnitudes of risk and for the neglect of some costs. While this is a sound argument, it does not provide an argument for other aspects of cost-benefit analysis. The second argument purports to show that commensurating and monetizing the values of the effects of regulation is warranted because it corrects for the use of widely diverging values of a statistical life. This argument fails because the use of widely diverging values of a statistical life is not a cognitive error: It is neither precluded by considerations of instrumental rationality, nor by the requirement of treating like cases alike.

Keywords: cost-benefit analysis, cognitive bias, risk, Cass Sunstein, ethics

1. Introduction

For[1] many of us, a moment's self-reflection would reveal that we are far from perfect when it comes to reasoning about probabilities, uncertainties and risk. In recent decades, cognitive psychologists, behavioural economists and others have corroborated this by providing evidence for the existence of a large number of systematic biases in people's thinking. Most of this work has had the descriptive aim of understanding how people in fact make judgments and decisions where risk is involved. But more recently, some have begun to draw explicitly normative conclusions from the psychological findings. One such conclusion is that society should regulate putatively risky activities on the basis of cost-benefit analysis, in order to avoid the negative influence of biases on risk regulation. The most prominent proponent of this 'cognitive argument' for cost-benefit analysis is Cass Sunstein. Over the last 15 years, Sunstein has articulated this argument in a number of writings (Sunstein, 2000; 2002; 2005; 2013, Ch. 7; 2014, Ch. 6-7). Focussing on Sunstein's account, I will argue that this cognitive argument fails, since the aim of avoiding cognitive error only gives us reason to instate one aspect of cost-benefit analysis – an aspect that is not specific to cost-benefit analysis nor necessitates the instatement of other aspects.

The plan of the paper is as follows: In section 2 I present Sunstein's argument and the evidence of cognitive biases he relies on. In section 3 I argue that we should make distinctions between different aspects of cost-benefit analysis as well as between different supposed cognitive errors. The upshot is that there are two cognitive arguments: One shows that the aim of correcting *false beliefs about the magnitude of risks* and *neglect of the costs of regulations* provide a reason for implementing the *analysis aspect* of cost-benefit analysis, and the other shows that the aim of correcting *divergences in the valuation of statistical lives* provide a

[1] I thank Sune Holm, members of the University of Copenhagen Analytic Philosophy Colloquium and two anonymous reviewers for helpful and quality-enhancing comments on earlier drafts of this paper.

reason for implementing the *commensuration* and *monetization* aspects of cost-benefit analysis. In section 4, I then argue that the former does not amount to an argument for cost-benefit analysis *tout court*, while in section 5 I argue that the latter argument fails because diverging valuations of statistical lives is not a cognitive error.

2. Cost-benefit analysis and psychological bias

The crux of Sunstein's argument is this: (i) A series of cognitive biases affect our thinking about risk; (ii) cost-benefit analysis (henceforth CBA) corrects for these; and (iii) this property of CBA gives us a reason to give CBA a very prominent role in the regulation of risk. In this section I will elaborate this claim by presenting in a little more detail the psychological claims Sunstein makes and the evidence that supports them.

First, however, it is necessary to briefly define CBA (I will discuss the content of CBA in more detail below). Generally speaking, CBA includes an analysis of the different positive and negative effects of some policy intervention, such as a regulation aimed to curtail some risk. A monetary value is assigned to each component effect, such as health effects, price increases, ecosystem degradation, and so on. These monetary values are drawn from the amounts people are willing to pay for the goods in question, or the amounts they require to be paid to accept some bad effect – ideally from market prices. Where no market prices exist, values are extrapolated from relevant market data. Where that is also impossible, survey studies are used. Common to all valuation methods in use is that they use values of constituent costs and benefits drawn from other contexts than the case under regulatory consideration.

Sunstein's cognitive argument stresses how CBA "corrects" cognitive biases. When discussing cognitive biases, we can distinguish between the psychological mechanism or tendency and the error to which it gives rise. The important thing for a cognitive argument

for CBA, of course, is that CBA corrects the *errors*. The various psychological mechanisms Sunstein discusses give rise to three types of error, namely false beliefs about the magnitude of risks, neglect of the costs of regulation, and divergences in the valuation of risks. While the psychological mechanisms involved may overlap and are interconnected in various ways, the three types of error they give rise to represent fundamentally different problems for our thinking about risk and its regulation.

2.1 False beliefs about the magnitude of risks

A risk can be understood as a combination of some *bad outcome* that might occur and the *likelihood* or probability that this outcome does occur. For example, a certain level of exposure to some substance might entail a 1-in-10,000 chance of developing a fatal cancer. Focussing, as Sunstein does, on the bad of premature death, the magnitude of a risk is equivalent to the likelihood of premature death, and, when summed over a population, to the expected annual number of fatalities. Sunstein suggests that several cognitive mechanisms lead people to form systematically wrong beliefs about the magnitude of various risks: People judge the likelihood of event-types (e.g. death from tuberculosis) on the basis of the cognitive availability of instances of the type (Sunstein, 2002, pp. 33-35; 2005, pp. 36-39); mechanisms of social transmission of information tend to amplify already erroneous beliefs about the magnitude of various risks; and various simplistic emotion-driven mechanisms influence judgment.

The result is a number of systematic errors in people's beliefs about the magnitude of risks. People overestimate the probability of low frequency events and underestimate the probability of high frequency events. They overestimate the number of fatalities from highly publicized and emotionally salient causes, such as accidents or cancer, and underestimate fatalities from "undramatic, quiet killers", such as diabetes (Liechtenstein et al., 1978). Risks

from activities that people view as highly beneficial, such as cell phone use, are underestimated, while risks from activities seen as being without benefits, such as GM agriculture, are overestimated. And when judging acts with a possible high-affect outcome, such as dying in a plane crash, the probability is sometimes neglected altogether.

The general effect of all of the mechanisms described above is that people's beliefs about the magnitude of risks caused by different activities and phenomena are systematically wrong. CBA corrects this by highlighting the actual magnitude of various risks, as well as by showing precisely how much risk-reduction a given policy would achieve.

2.2 Neglecting the costs of regulation

The activities that produce risks typically also have benefits. According to Sunstein, ordinary people tend to neglect these benefits when they evaluate risky activities. Since the costs of regulation are partly the mirror image of the benefits of the regulated activity, this means that people tend to neglect some *costs* of regulation. One cause of cost neglect is people's tendency to perceive risky activities as having few benefits (and hence regulating them to have few costs). Another cause is loss aversion. Since people dislike losses more than they like corresponding gains, they tend to focus on possible losses – i.e. risks – and neglect the benefits that are foregone by regulating. Furthermore, there is a tendency to judge losses and gains relative to the status quo, leading to a higher tolerance for existing risks than new ones. Finally, the systemic side effects of trying to remove a risk are often ignored. In many cases, regulations will create "substitute risks" – for example, banning asbestos in car brakes may increase the risk of brake failure. Such "risk-risk tradeoffs" are often ignored (Sunstein, 2002, pp. 39-40; 2005, pp. 32-33).

For these reasons, Sunstein argues, there is a tendency for people to ignore the costs of regulatory policies. As he puts it, the benefits of regulation are "on-screen" while the

costs are "off-screen". Some studies suggest that *merely* placing costs on-screen leads people to alter their judgments (Sunstein 2002, p. 42; 2005, p. 48). CBA corrects cost neglect because all effects of regulation, costs as well as benefits, are present on CBA's screen.

2.3 Diverging valuation of risks

The final cognitive error that Sunstein identifies is the assignment of (widely) diverging monetary values to reducing risks that are statistically expected to claim the same number of lives – or, equivalently, assigning (widely) diverging monetary values to a 'statistical life'. Sunstein clearly takes diverging valuations to be a major problem: In the two places where he deals exclusively with the cognitive argument for CBA, he begins by discussing a study showing that the cost per premature death averted for a large set of regulations varied from $0.1 million to $92 billion (Sunstein, 2000, pp. 1061-1064; Sunstein, 2002, pp. 29-33). The fact of divergence is thus the primary datum that motivates a closer scrutiny of risk regulation practices. Furthermore, Sunstein describes the divergences uncovered as "anomalies" and claims that they are, at least in part, caused by irrationalities (Sunstein, 2002, pp. 48-49). In addition, Sunstein suggests that (widely) diverging valuations of statistical lives are marks of a system of regulation that is not "coherent" (Sunstein, 2005, p. 149).

Diverging valuations is less obviously a cognitive error than the two categories above. Sunstein gives two reasons for believing that diverging valuations are nevertheless problematic. First, the assignment of diverging values to saving a statistical life "creates a presumption that [a] system of regulation suffers from serious misallocation of resources" (Sunstein, 2002, p. 31). In other words, a regulatory system that assigns widely diverging monetary values to statistical lives is using available resources unwisely, and could do more good by reallocating resources to make the valuations of statistical lives roughly the same across cases. Sunstein cites a study that suggests that reallocating resources in this manner

could save an additional 60.000 lives annually (holding cost constant), or save $31 billion, (holding lives saved constant) (Sunstein, 2002, p. 25; the study cited is Tengs & Graham, 1996). Second, he appeals to the uncontroversial principle that like cases should be treated alike to argue that the value of reducing one risk should only be different from the value of reducing another risk if there is a morally salient difference from one case to another (Sunstein, 2002, p. 48; 2005, p. 149; Sunstein, Kahneman, Shkade & Ritov, 2002, p. 1153).

All of the abovementioned psychological effects can generate diverging valuations. In addition, divergences are caused by two further phenomena. The first is what Sunstein calls "the proportionality effect", which denotes a preference for saving a larger proportion of people in the reference class rather than saving a larger absolute number (Sunstein, 2002, pp. 47-48). For example, many would prefer to save the people living in a village of 100 from a 1-in-20 risk of death to saving the people of a city of 10,000,000 from a 1-in-1,000,000 risk of death, although the latter act is expected to save more people (10 rather than 5 statistical lives). Consequently the amount of resources devoted per statistical life saved is likely to be higher for small populations facing a somewhat higher probability of death than for larger populations facing a somewhat lower probability of death. The second is "separate evaluation", which occurs when one policy is evaluated without reference to other policies. In a study, Sunstein and co-authors found that people's judgments concerning the value of providing two different benefits changed markedly when these benefits were assessed jointly rather than separately (Sunstein, Kahneman, Shkade & Ritov, 2002, pp. 1174-1178). Such "joint evaluation" is, of course, characteristic of the coherentist methods that are dominant within moral epistemology.

According to Sunstein, CBA provides a corrective for the diverging valuations problem in three ways. First, it provides information about how many lives are saved and at what cost, and thereby prevents divergence from emerging from unreflective processes.

Second, it directly prevents separate evaluation by using evaluation methods that make references to other cases than the one being treated. Third, CBA directly limits divergence by institutionalizing a presumptive value for a statistical life that holds across cases.

3. Evaluating the cognitive argument

Just as we can distinguish between three types of cognitive error, we can distinguish between three aspects of CBA, which I dub *analysis, commensuration* and *monetization*. Drawing these distinctions reveals a more complex picture behind the claim that CBA corrects cognitive errors: We should ask *which aspects* of CBA corrects for *which errors*. The three aspects of CBA are as follows:

Analysis is an assessment of the likely effects of the various available policy alternatives – how many lives would be saved, what other health problems would be avoided, which ecosystems would be preserved, how much would it cost industry to implement, etc. Analysis provides an accounting of the effects using measures that are 'natural' for the given kind of effect.

Commensuration is the assignment of values to the various effects that allow a comparison of the values of different things (e.g. lives saved and industry costs) – or, equivalently, that allow us to talk of the *relative* values of these different things. Commensuration is necessary for generating meaningful sums of benefits and costs. Importantly, the type of commensuration that is part of CBA assigns values to *constituent* costs and benefits that are supposed to apply *across* cases.

Monetization denotes the use of money as the metric that allows commensuration. The standard method for generating monetized values is by estimating ordinary people's willingness-to-pay (WTP) or willingness-to-accept (WTA). That is, the analyst estimates how much people are willing to pay in order to avoid, for example, a certain risk of death or the

destruction of an ecosystem (WTP), or alternatively how much they would need to be paid in order to allow themselves to be subjected to a certain risk, or to allow an ecosystem to be destroyed (WTA). Where the thing in question is not directly traded in a market, WTP/WTA is extrapolated from relevant market behaviour – especially wage premiums for risky occupations where mortality risks are concerned – or estimated on the basis of surveys.

Now, in virtue of which of these aspects does CBA correct which of the cognitive errors? Consider first false beliefs about the magnitude of risks. This error is solved through analysis, since that aspect provides a full assessment of the likely effects of regulation. This includes risk information in terms of the probabilities of adverse effects, the number of people bearing each type of risk, and/or the aggregated expected number of deaths caused or prevented by regulation. It is not necessary to translate risk numbers into any other value in order to get them right; arguably, such translation is more likely to distort the information. Second, consider neglect of the costs of regulation. Again, this error is prevented by analysis, which includes an assessment of all effects of regulation, including costs. It is not necessary to commensurate costs with benefits in order to bring them "on-screen". Finally, consider diverging valuations of risks. Here, analysis is not sufficient, since it does not even provide numbers that allow comparisons of diverging valuations across cases. Hence commensuration and monetization are needed to avoid diverging valuations.

The upshot is that there really are two cognitive arguments: The first argument says that false beliefs and cost neglect are cognitive errors, and that the fact that analysis corrects those errors provides a reason for using analysis. The second argument says that diverging valuations is a cognitive error, and that the fact that commensuration and monetization corrects that error provides a reason for using commensuration and monetization. In the next two sections, I will argue that neither of these arguments succeeds in providing a cogent cognitive argument for CBA.

4. Is the argument for analysis an argument for CBA?

Why might one think that an argument for *one aspect* of CBA counts as an argument for CBA *as a whole*? I can think of two reasons. First, one might simply identify CBA with analysis. This would represent a divergence from what is standardly meant by CBA (although not one that is *entirely* unheard of (Adler & Posner, 2006, p. 79)). If one were to make an argument of this sort, one ought to flag strongly that the sense of 'cost-benefit analysis' used is non-standard, and avoid intermingling the argument with discussions of commensuration and monetization. Sunstein does neither of these things. However, he does frequently describe CBA as being a mere account of all effects of the regulations considered, and as no more than an informational input that does not constrain decisions, both of which sound more like analysis only than full CBA (e.g. Sunstein, 2002, pp. 35 & 106-110; Sunstein, 2005, p. 129).

If this is what Sunstein means by 'the cognitive argument for CBA', that argument is of course sound. But I do not think this is in fact what Sunstein means. The conclusion he draws from considering cognitive biases is that there should be an "incompletely theorized agreement" among proponents of all or most plausible normative outlooks on eight concrete proposals to be implemented in legal and administrative practice (Sunstein, 2002, p. 110-113). The implementation of these eight proposals would, taken together, amount to the implementation of all three aspects of CBA, including commensuration and monetization. Under Sunstein's scheme, agencies charged with risk regulation would be required to show that the benefits outweigh the costs before being allowed to implement the regulation. Furthermore, agencies would be required to monetize all effects and to use a value of a statistical life (henceforth VSL) that falls within a set range (Sunstein, 2002, p. 111-112; 2014, ch. 2). The current recommended range is a floor of $1 million and a ceiling of $10 million (Sun-

stein, 2014, p. 204, n. 12). This interval is based on studies intended to measure people's willingness-to-pay/willingness-to-accept – i.e. they are the standard CBA value measures (Sunstein, 2005, ch. 6; 2014, p. 51). In conjunction, these requirements entail that a regulation is only acceptable if its benefits outweigh its costs, where costs and benefits are monetized in the manner characteristic of CBA. In other words, they implement *full* CBA in a way that constrains decisions.

However, it should be noted that the CBA advocated by Sunstein is a soft, non-rigid form. Both the requirement to show that benefits justify costs and the VSL interval are presumptive only. It would be possible for regulative agencies to go ahead with regulations whose costs exceed their benefits "on the basis of a publicly articulated explanation", e.g. in court, or in response to regulatory review (Sunstein, 2002, p. 112). Agencies would likewise be allowed to use a VSL outside the set interval if they can provide a reasoned explanation. In particular, agencies should be allowed to adjust the VSL in accordance with various "qualitative factors". These presumptions make it unclear what the effects of Sunstein's proposals would be. Much would be left to the discretion of judges and civil servants. Nevertheless, his proposals would certainly do more than merely implement analysis.

The second possible reason for thinking that the cognitive argument for *analysis* is a cognitive argument for *CBA* is to take the fact that one aspect of CBA possesses a reason-giving feature to provide a reason for implementing CBA as a whole. The general principle is not unfamiliar: The durability of the engine provides a reason to buy the entire car, and the deliciousness of the dessert provides a reason to buy the whole tasting menu. Since the analysis aspect possesses the reason-giving property of correcting false belief and cost neglect, it would not be strange to say that there is a (cognitive) reason for using CBA. However, it *would* be strange to claim that the properties of one aspect of a whole provide reason for favouring other aspects of that whole, e.g. that the engine's durability counts in favour

of the car's leather interior, or that the fact that the tasting menu's dessert is delicious provides a reason to buy that menu's entree from the à la carte selection. Similarly the desirable qualities of analysis provide no reason for using commensuration or monetization.

There is thus some basis for saying that the cognitive argument for analysis is a cognitive argument for CBA. However, it is a limited argument, in two ways. First, it does not provide a reason for implementing commensuration and monetization when we are considering whether to implement aspects individually. In the case of Sunstein's eight propositions, there seems to me no good reason to suppose that they cannot be implemented separately. Consequently, no incompletely theorized agreement can be expected on those propositions that implement commensuration and monetization.

Second, if the cognitive argument for CBA is at bottom an argument for analysis only, then no reason has been given for *preferring* CBA to any alternative that also includes analysis. Consequently we have been given no reason for preferring CBA to alternatives that already include analysis, or for merely reforming alternatives that do not include analysis but are compatible with it. Many of CBA's primary rivals already include analysis. Consider, for example, what Adler & Posner (2006, p. 73) call the "intuitive balancing" approach, which they claim is widely used. Under this approach "[p]olicy effects will be described, and indeed might be quantified on various scales (for example, numbers of deaths, acres of ecosystem destroyed, jobs lost or gained)". This is exactly what analysis is. However, the intuitive balancing approach does not commensurate or monetize, but weighs the costs and benefits in given case in an intuitive manner.

Another widely used approach is technology-based regulation (Driesen, 2005; 2011; McGarity, 2002, p. 2343-2344). Under this approach, agencies operate with two aims: One is an ideal but non-enforceable state wherein no harm is expected to occur; for example, a level of exposure to a toxin that is known not to be harmful. The second is a non-ideal, but

enforceable goal, that requires risk producers the move as close to the ideal goal as is technologically feasible; for example, to reduce emissions of the toxin so as to bring exposure as close to the safe level as possible. 'Feasibility' here means both that the given limit to risk must literally be technically possible and that it must be economically possible for the industry regulated. An economically impossible regulation is typically taken to be a regulation that makes it impossible for large sections of the relevant industry to continue operating (Driesen, 2005, p. 9). Technology-based regulation thus includes an analysis of risks – required for setting the ideal goal – as well as an analysis of costs. But it neither commensurates nor monetizes effects.

I conclude, therefore, that the route to a successful cognitive argument for CBA through the (assumed) successful cognitive argument for analysis fails, since it does not provide a reason for implementing CBA as a whole where analysis is available on its own, and since it does not provide a reason to prefer CBA to most important rivals. It does, of course, provide a reason to prefer CBA to alternatives that do not include analysis, such as directly and uncritically responding to public demands for regulation. But it is misleading call this a cognitive argument *for CBA* when it is in fact equally an argument for a number of plausible alternatives.

5. Diverging valuation and cognitive error

Although the route to a cognitive argument for CBA through the argument for analysis fails, the route through the argument against diverging valuations may yet succeed. Following Sunstein, I will only discuss the valuation of statistical lives. If diverging VSLs are a cognitive error, and if commensuration and monetization corrects that error, then there is a cognitive argument for commensuration and monetization. In conjunction with the cognitive argument for analysis, this would amount to a cognitive argument for CBA as a whole.

Although diverging VSLs is less obviously a cognitive error than false beliefs and cost neglect, I suggested earlier that there are two kinds of reason for thinking that it is: First, that diverging VSLs amount to a suboptimal allocation of resources, and second, that using diverging VSLs means treating similar cases differently. Call the first *instrumental* reasons and the second *coherence* reasons. I will now argue that neither of these types of reason succeeds in showing that diverging valuation is a cognitive error that can be corrected by commensuration and monetization.

5.1 Instrumental reasons

Assume that all we care about is saving as many lives as possible. Then any approach that results in fewer lives than possible being saved is instrumentally irrational. Sunstein argues that current (or rather pre-CBA era) regulatory policy suffers from exactly that instrumental irrationality. As mentioned, he frequently refers to a table showing widely diverging VSLs, and suggests that large benefits could be realized by reallocation. The numbers in that table are somewhat old, not to mention controversial (Heinzerling 1998; Parker 2003). As a more recent figure, Sunstein cites calculations by the Office of Management and Budget (OMB) showing an increase in net benefits from regulation from $3.4 billion annually to $91.3 billion during the first three years of the Obama administration, when CBA was widely implemented under Sunstein's own direction (OMB, 2012, p. 59; Sunstein, 2013, pp. 33-35).

These claims of waste before and improvement after CBA was used can certainly be questioned. With respect to the latter, there is a measure of question begging involved, since the calculations made by the OMB are based on monetized costs and benefits as recommended in CBA. It is not surprising that a policy aimed at maximizing CBA-defined net benefits in fact increase CBA-defined net benefits. But this cannot be used as a *neutral* measure of how much good is done by using CBA.

With respect to the former, the mere fact that some regulations spend smaller sums per statistical life saved than others is not sufficient to show that more lives could be saved. First of all, what matters is not actual expenditure per life saved, but the marginal cost of saving an *extra* life for each regulation. There is no general reason to think that actual expenditure is a good indicator of the marginal cost of an extra life saved: Some regulations may have succeeded in eliminating mortality risk altogether, leaving no extra lives available to be saved; or the potential of one relatively cheap method of reducing risk may have been reached, and any alternative method, e.g. banning the activity altogether, may be much more expensive.[2] Second of all, the savings from deregulating cannot always be reallocated to other, cheaper-per-life-saved regulations. In many cases, the savings associated with deregulating accrue to various private actors (e.g. businesses, workers or consumers). Such savings are not immediately available for reallocation, although they could in principle be made available, e.g. by taxing risk production.

Suppose, however, that we could in fact save more lives at the same cost by reallocating our regulatory efforts. This fact alone still does not warrant the use of a full CBA that monetizes the value of lives saved. Suppose we were to hold current expenditure fixed and then go about saving lives in the most cost-efficient way. In that case, the VSL would be determined by what the cost of saving the final life we could afford to save is. There would be no need to determine the VSL by any external method, such as WTP/WTA.

The effect of using an externally determined VSL is to implicitly set our total risk reduction 'budget' – the total expenditure we would incur by saving all lives that could be saved at less than $X each. There is an allocation-based rationale for using the externally

[2] The study by Tengs and Graham cited by Sunstein does take the marginal cost issue into account by using numbers for the "level of implementation" of various interventions. However, those numbers are estimates by anonymous reviewers, and there is no publically available information about what their basis is (Parker 2003: 1362-1363, 1377-1381).

determined VSL, but it has to do with what *other* goods we could get instead of lives saved, rather than with how many lives we could save at a given cost. Only if the externally determined VSL captures a value ratio between risk reduction and everything else that is (in some sense) *correct* is it instrumentally irrational to spend more than the externally determined VSL on saving a statistical life (since we are getting less overall good than we could have gotten via an alternative allocation of resources).

5.2 Coherence reasons

Even if diverging VSLs is not evidence for misallocation, they may still be indefensible for coherence reasons. Sunstein argues that diverging VSLs are evidence of a lack of coherence in our regulatory system (Sunstein, 2005, p. 149). In other words, they are evidence that we are arbitrarily treating cases differently. He also suggests that CBA avoids such incoherence because it implements the coherentist method: Since the VSL used in CBA that is imported from a different case than the one under consideration, it is a form of joint, rather than separate, evaluation. *Prima facie*, this idea is plausible. CBA is at least not guilty of conducting a separate evaluation *of the regulatory case under scrutiny*. But no *joint* evaluation is at work either. Rather, the VSL is typically extracted from a single (type of) case, namely the wage premiums workers receive for taking jobs that include mortality risks. The 'comparison' between wage premium cases and regulation cases is entirely one way – no argument is ever made that wage premiums are too high, or too low, based on feedback from other cases. CBA thus substitutes separate evaluation of one case for separate evaluation of another case. It is not a coherentist method.

Although CBA is not a coherentist method, we might still suspect that diverging VSLs reveal arbitrarily differential treatment of cases. This idea, too, has some *prima facie* appeal. If we are spending $1 million to save a life in the domain of workplace accidents but

$20 million in the domain of air traffic, it seems that we are valuing the lives of workers at only one twentieth of the lives of air travellers. But this is only true if the VSL in some sense measures the moral importance we are attaching to each person's life. It need not do so. There is an ambiguity in the expression 'the value of a statistical life'. In one sense, the VSL refers to an *input* into a decision procedure that assigns a value to constituent effects of regulation and makes a decision on the basis of the balance of values (i.e. a decision procedure that commensurates). Call this an *input VSL.* In the other sense, the VSL is simply an accounting consequence of a regulation that is based on a rationale that does not employ a (monetized) input VSL: If the cost of the regulation is $100 million and it saves 10 statistical lives, then the VSL of that regulation is $10 million. Call this a *revealed VSL*. Only where the decision procedure employs a monetary input VSL can the monetary VSL be said to measure the moral importance being placed on the statistical lives that are at stake. Thus it is only in those cases that diverging VSLs can be said to constitute diverging moral valuations of lives. But not all reasonable rationales for regulation use a monetized input VSL. They may use a non-monetized input VSL, or use inputs other than VSL, or not use inputs at all.

The way Sunstein sets up his argument blurs the distinction between input VSL and revealed VSL. He proposes that regulatory agencies should be *required* to produce a CBA. In other words, they are required to produce an argument for regulation framed in the terms of the regulatory rationale of CBA – i.e. a rationale that commensurates, takes VSL as one of the inputs, and that monetizes this input VSL. If this is a requirement that only a CBA-style regulatory rationale is eligible, Sunstein's argument is viciously circular. But suppose agencies were allowed to base their decision on another type of rationale – one that uses non-monetized input VSL, or uses inputs other than VSL, or does not use inputs at all – but were still required to produce a CBA inclusive of a monetary VSL. In that case, the monetary VSL 'used' by agencies is a revealed VSL, not an input VSL. Hence the fact that agencies use

diverging monetary VSLs do not give us reason to think that they are assigning different moral importance to different people's lives.

The claim that diverging VSLs constitute a cognitive error is, therefore, unwarranted if non-CBA rationales for regulation are reasonable. Below I will argue that they are. Although I will perhaps also be insinuating that they are *more* reasonable than the CBA rationale, this is not a necessary part of my argument against the cognitive argument for CBA. Since the cognitive argument for CBA requires that there could be an incompletely theorized agreement on a legally binding (although presumptive) interval for the VSL, it is sufficient to show that plausible standpoints do not agree that there is a problem to be solved at all, and/or that a mandated interval of monetized VSLs is the solution to what problems might exist.

5.2.1 No inputs

Some rationales for regulation do not use inputs; i.e. they do not assign values to the constituent effects of regulation that are assumed to be applicable across all cases. On the intuitive balancing approach, for example, regulators judge whether the benefits of a proposed regulation taken as a whole outweigh the costs taken as a whole. They thus do not necessarily assign a value to constituent costs, and the weighing will be case-specific. And in technology-based regulation, no inputs are used at all, since regulation is simply determined by what the highest feasible level of protection from risks is.

These rationales for regulation are rationalized by forms of non-consequentialist theory. Many non-consequentialists argue that aggregation – i.e. weighing costs against benefits – is not permitted in cases where the members of one group stand to get a minor benefit while the members of the other group stand to get a major benefit (e.g. Scanlon, 1998, pp. 235-236). Now, suppose group A will bear the costs of a regulation while group B will bene-

fit (by being relieved of bearing a risk). If the cost to each member of A is trivial compared with the risk that the members of B will otherwise bear, then on a non-consequentialist view we should help B – that is, implement the regulation – even if the aggregate benefits to group A are larger (see Lenman, 2008; James, 2012; Frick, 2015). These non-consequentialists thus focus on the *trade-off* made in a given case instead of the values of each constituent effect considered independently.

From the point of view of non-consequentialist theories of this kind, both intuitive balancing and technology-based regulation make a lot of sense. Since the trade-offs made in each case are what matters, any balancing of costs and benefits should be expected to be case-specific, as it is in intuitive balancing. And technology-based reasoning can be viewed as the implementation of a specific set of allowable trade-offs: Predictable deaths are not to be traded off against mere economic loss to a risk-producing industry, but predictable deaths *can* be traded off against the costs associated with undermining the industry's ability to operate.

5.2.2 Non-VSL inputs

Even if we accept that a commensurating decision procedure is the best one, it is an open question what inputs we should use in such a procedure – or, in other words, what the morally *important* constituent effects of regulation are. If the inputs used by regulators are not VSLs, then the revealed VSL of regulation may vary. And a number of plausible positions in ethics imply that the VSL is not the relevant input. For example, on the standard account, the badness of a death is proportional to the life span lost (see e.g. McMahan, 2002, ch. 2; Nagel, 1970). For those holding such a view, the relevant input to decision-making should not be statistical *lives* lost, but statistical *life years* lost. Consequently, the revealed VSL of interventions that save younger people should (*ceteris paribus*) be higher than the

revealed VSL of interventions that save old people. Another example is the type of death involved. For instance, it is plausible to hold that a death resulting from a long and painful illness, such as cancer, is worse than instantaneous death from a workplace accident. If regulators use *types* of death as input, the revealed VSL for interventions that prevent diseases like cancer will be higher than the revealed VSL for interventions that prevent workplace accidents.

Sunstein is well attuned to this problem, and in fact he believes that the VSL is the *not* important input (Sunstein, 2004; 2014, ch. 4). Instead, the input he recommends is the "values ordinary people assign to risks" – i.e., their WTP/WTA (Sunstein, 2005, p. 131). This raises the question of why Sunstein believes that widely diverging VSLs is a cognitive error that needs to be controlled by a legally mandated VSL interval. The answer is that it is because *WTP* falls within such an interval. Sunstein claims that the VSL interval currently used by the U.S. government is congruent with "the range suggested by the current technical literature", where this technical literature consists of WTP studies (Sunstein, 2014, p. 51). Furthermore, the cases in which he envisages the presumptive interval to be overridden are cases where studies suggest WTPs are markedly different, e.g. in the case of cancer deaths (Sunstein, 2005, pp. 139-141).

However, not all reasonable views would agree with Sunstein on this. First of all, the view that WTP is the correct input is controversial. Second, and more importantly, Sunstein's assumption that WTP falls within a relatively narrow interval relies on the rejection of some of the variation in WTP that have been observed (Sunstein, 2002, Ch. 3; 2005, Ch. 7). In particular, he disregards two important differences between risk cases, namely *voluntariness* and the *distribution* of risks and the benefits from the risk's existence. These factors alter ordinary people's assessments of risk drastically (Slovic, 2000). Furthermore, and more im-

portantly, giving proper weight to these factors cast doubt on the justifiability of using a relatively narrow VSL interval on moral grounds.

Consider first voluntariness. When considering the different values of a risk that is borne voluntarily and one that is not we must consider the importance of a number of phenomena, including freedom, consent, compensation, and personal responsibility. Arguably, the proper moral role of these factors is so complex and disputed that it defies a framing in terms of inputs at all.[3] But supposing that we can frame the issue in terms of inputs, there is no reason to expect the value of a voluntary risk of death (e.g. from skydiving) and a non-voluntary risk of death (e.g. from air pollution) to be even roughly the same, even if the probabilities of death are the same.

Second, consider distributions of the risks and benefits from the risk's existence (see Hermansson & Hansson, 2007). There are two very simple kinds of case and a range of unsimple ones. In one kind of simple case, the beneficiaries of regulation also bear the entire cost. Call these Type A cases. In the other, the opposite is true; the beneficiaries and the cost bearers are wholly distinct groups. Call these Type B cases. In between are combinations of the two, where different persons incur some of the cost and/or some of the benefit of a regulation. For simplicity, I focus only on the simple types of case. Type A and Type B cases involve very different moral judgments. In Type A cases, autonomy will play a large role, since we are deciding on behalf of the individual what her trade-off between bearing the risk and enjoying the benefits of the risk-producing activity should be. In Type B cases, the judgment may be about how much money we, as a society, should be willing to spend to reduce risks, or it may be about what burdens we may legitimately place on *others* in order

[3] As Sunstein rightly argues (2002, Ch. 3), voluntariness is not plausibly conceived of as an either/or property, but rather reflect a number of underlying factors that differ in degree.

to reduce a risk to a given set of persons. There is no *prima facie* reason to think that these very different judgments should yield similar results.

So giving appropriate moral weight voluntariness and distribution seems to militate against a narrow VSL interval. What is more, it suggests that WTP numbers systematically undervalues VSL in a large number of cases. First of all, WTP studies normally concern Type A cases. In Type A cases, the rationale for using WTP, rather than a higher number, is the plausible principle that we should not force people to make exchanges that they would rather have avoided (Sunstein, 2005, pp. 150-153). But this rationale is not applicable at all in Type B cases, where no individual is forced to make any exchange. Second, if WTP studies must assume that (for example) wage differentials reflect the well-informed and autonomous valuation workers place on risks, which implies that the risks are voluntarily borne. If this assumption is true, the reason for using WTP to determine the VSL in cases of involuntarily borne risk is undermined. And, assuming that involuntarily borne risks are worse than voluntarily borne ones, CBA systematically undervalues statistical lives lost due to involuntarily borne risks.

5.2.3 Non-monetary VSL

It is perhaps not surprising that divergence in revealed VSL follows from regulatory rationales that do not use inputs or that use inputs other than VSL. But even if we assume that the VSL is the same across all cases, the *monetary* revealed VSL may still be different if the regulatory rationale uses a non-monetized measure of the VSL. Although monetization is frequently presented as a mere pragmatic tool by CBA proponents, no alternative has been tried in practice (as far as I know). But consider as a contrast an approach which assigns to each constituent effect a set number of 'regulatory points', which are supposed to measures the moral importance of each constituent effect. Suppose also that this 'points approach'

assigns the same number of points to each statistical life. The revealed VSL of a regulation designed on the basis of the points approach would depend on how many points the various *costs* of regulation are assigned. For example, if the moral importance of a $100 million price increase has a different moral importance than $100 loss to shareholders, then these costs would be assigned different points – and the revealed *monetary* VSLs of regulations saving equal numbers of lives per regulatory cost point would consequently differ depending on the type of cost involved.

It is very plausible that not all costs are equally morally important. For example, those who bear the cost may have no claim to the good of which they will be deprived. This is especially true when those who bear the costs of regulation are also causally responsible for the existence of the risk and do not bear the risk themselves – that is, when they *impose* the risk on others for their own benefit. But more generally, there seem to me to be little reason to assume that a $100 million increase in taxes on leaded gasoline has the same moral importance as $100 million of lost profits to farmers from a ban on a cheap but toxic pesticide, and that both of these have the same moral importance as a $100 million worsening of the government budget balance.

A further problem is that the use of monetary values as inputs in CBA means that effects that are 'naturally' measured in money are not subjected to the same scrutiny as effects that are not naturally counted in money, such as adverse health effects and environmental degradation. Consider, for example, Sunstein's description of how 36 ideal-type cases of regulation would be (and presumably were) examined using CBA during the process of regulatory review that he administered (Sunstein, 2014, Ch. 2). No questions are posed on the cost side – in every case the cost is merely stated as a total aggregate dollar amount. By contrast, the benefits of regulation are meticulously scrutinized. At least one reason why CBA's critics are often especially averse to the *monetization* aspect of CBA is exactly that it

makes the analysis blind to morally important differences on the cost side, where the 'natural' currency is likely to be money. In other words, the monetization step will very likely have the effect that morally relevant differences between costs are erased. The fact that CBA fails to rationally scrutinize one side of the cost-benefit equation is especially troubling for an argument that conceives of CBA as an engine for rational scrutiny, as the cognitive argument does.

6. Conclusion

I have argued that Sunstein's cognitive argument fails. Closer inspection reveals that both 'cognitive error' and 'cost-benefit analysis' are concepts that can be decomposed into different phenomena, and that the cognitive argument for CBA is better viewed as two separate arguments. One argument is persuasive, but is not an argument for CBA, only for the analysis aspect of it. The other argument would provide reason to implement CBA as a whole, but is unpersuasive since we have neither instrumental nor coherence-based reasons to treat diverging valuations of statistical lives as a cognitive error. Instrumental reasons rely on the false (or at least unsubstantiated) assumption that reallocating resources as CBA recommends would realize large gains, and they fail to justify setting the total risk-reduction budget at any given level. Coherence-based arguments falsely assume either that CBA is a coherentist method or that monetary VSLs are meaningful measures of the moral importance attached to saving a life, even where monetary VSLs are merely revealed rather than used as inputs. Consequently, there is no cogent cognitive argument for cost-benefit analysis in risk regulation.

References

Driesen, D.M. (2005). Distributing the Costs and Benefits of Environmental, Health and Safety Protection: The Feasibility Principle, Cost-Benefit Analysis and Regulatory Reform, *Environmental Affairs, 31*(1), 1-95.

Driesen, D.M. (2011). Two Cheers for Feasible Regulation: A Modest Response to Masur and Posner, *Harvard Environmental Law Review, 35*(2), 313-341.

Frick, J. (2015). Contractualism and Social Risk, *Philosophy and Public Affairs*, *43*(3), 175-223.

Heinzerling, L. (1998). Regulatory Costs of Mythic Proportions, *The Yale Law Journal*, *107*(7), 1981-2070.

Hermansson, H. & Hansson, S.O. (2007). A Three-Party Model Tool for Ethical Risk Analysis, *Risk Management*, *9*(3), 129-144.

James, A. (2012). Contractualism's (Not So) Slippery Slope, *Legal Theory*, *18*(3), 263-292.

Lenman, J. (2008). Contractualism and Risk Imposition, *Politics, Philosophy and Economics*, *7*(1), 99-122.

McGarity, T.O. (2002). Professor Sunstein's Fuzzy Math, *The Georgetown Law Journal*, *90*, 2341-2377.

McMahan, J. (2002). *The Ethics of Killing: Problems at the Margin of Life.* Oxford: Oxford University Press.

Nagel, T. (1970). Death, *Noûs*, *4*(1), 73-80.

OMB (2012). *2012 Report to Congress on the Benefits and Costs of Federal Regulations and Unfunded Mandates on State, Local, and Tribal Entities*. Retrieved from https://www.whitehouse.gov/sites/default/files/omb/inforeg/2012_cb/2012_cost_benefit_report.pdf

Parker, R.W. (2003). Grading the Government, *The University of Chicago Law Review*, *70*(4), 1345-1486.

Slovic, P. (2000). *The Perception of Risk.* London: Earthscan.

Sunstein, C.R. (2000). Cognition and Cost-Benefit Analysis, *The Journal of Legal Studies, 29*(52), 1059-1103.

Sunstein, C.R. (2002). *Risk and Reason: Safety, Law and the Environment*. Cambridge: Cambridge University Press.

Sunstein, C.R. (2004). Valuing Life: A Plea for Disaggregation, *Duke Law Journal, 54*(2), 385-445.

Sunstein, C.R. (2005). *Laws of Fear: Beyond the Precautionary Principle*. Cambridge: Cambridge University Press.

Sunstein, C.R. (2013). *Simpler: The Future of Government*. New York: Simon & Schuster.

Sunstein, C.R. (2014). *Valuing Life: Humanizing the Regulatory State.* Chicago: The University of Chicago Press.

Sunstein, C.R., Kahneman, D., Schkade, D. & Ritov, I. (2002). Predictably Incoherent Judgments, *Stanford Law Review*, *54*(7), 1153-1215.

Tengs, T.O. & Graham, J.D. (1996). The Opportunity Cost of Haphazard Social Investments in Life-Saving. In R.W. Hahn (Ed.), *Risks, Costs and Lives Saved: Getting Better Results from Regulation*. Oxford: Oxford University Press.